contact@namiq.io · namiq.io

NamiQ

SECURITY ADVISORY

Project MONAI MONAI Label exploited

CVE-2026-133968 · on CISA's exploited list since 23 November 2026

Analysis Report

Anyone running MONAI Label 0.8.5 or earlier must upgrade to 0.8.6 by 26 November 2026 — CISA's deadline is tomorrow. Authentication is off by default, making every reachable instance exposed; check …

PREPARED BY	NamiQ	DATE	25 November 2026
FACTS AS OF	25 November 2026	VERSION	1.0
SEVERITY	**Critical: CVSS 3.1 9.8**	STATUS	On CISA's exploited list; due 26 November 2026

NamiQ © 2026 NamiQ. All rights reserved. This document is the property of NamiQ. It may be shared unaltered, with attribution.
Companion article: namiq.io/insights/project-monai-monai-label-cve-2026-133968

The bottom line

> *Anyone running MONAI Label 0.8.5 or earlier must upgrade to 0.8.6 by 26 November 2026 — CISA's deadline is tomorrow. Authentication is off by default, making every reachable instance exposed; check whether MONAI Label is installed anywhere on your network first.*

At a glance

Fact	Detail
Product	Project MONAI MONAI Label (CVE-2026-133968)
Severity	Critical: 9.8 (CVSS 3.1)
CISA	On its exploited list since 23 November 2026. Deadline for US federal agencies: 26 November 2026.
Chance of exploitation (EPSS)	0.27, the estimated chance in the next 30 days, as at 25 November 2026
Ransomware use	Not recorded by CISA. That is not the same as none.
Fixed in	0.8.6, released 22 November 2026. The release notes name both a pip package and a Docker image at that version.
Who looks first	Teams running a MONAI Label server. Authentication is off by default, so any client that can reach the server is a potential attacker. The vendor's advisory says that with authentication on, the ANNOTATOR role is sufficient to exploit this.

Does this apply to you?

Answer in order. Stop at the first no.

Question	If no	If yes
Do you run a MONAI Label server anywhere on your network?	Nothing to do for this CVE.	Check the version running (next question).
Is it running version 0.8.5 or earlier?	Not in the affected range; no action needed for this CVE.	Treat as urgent; continue.
Is it reachable from outside your internal network, or is authentication disabled?	Lower urgency, but patch before the next maintenance window.	High priority — patch or restrict access immediately.

What it is

An attacker sends a crafted image identifier to the server's upload endpoint. The server constructs a file path from that identifier without checking it stays inside the datastore, then writes the uploaded file there. The vendor's advisory says this can lead to remote code execution or denial of service.

Why it matters

The record says this has been exploited in the wild since September 2026. Authentication is off by default, so in a standard deployment any client that can connect is enough. A source reports that attackers identified the vulnerable endpoint and exploited it in approximately 11 minutes. Our reading is that any server running 0.8.5 or earlier that was reachable from an untrusted network should be treated as potentially compromised and reviewed for unexpected file changes before returning to service.

What the vendor's advisory says

We read three pages on github.com and www.theregister.com. In short:

- A second path traversal is also fixed in 0.8.6: the session upload endpoint extracted tar archives without checking member names, and session IDs were used as paths without validation. Affects all versions up to and including 0.8.5.
- Workarounds until upgrade: enable the MONAI Label authentication setting, restrict server access to trusted networks, and run the server as an unprivileged user with write access limited to the datastore.
- The release notes say any 0.8.5 or older server reachable from an untrusted network should be reviewed for unexpected file changes before returning to service.

Things people will get wrong

- Enabling authentication fully protects us from this flaw. The vendor's advisory says the ANNOTATOR role is sufficient to exploit this. Enabling authentication reduces risk only if that role is not granted to untrusted users.
- We only use this internally so we are not at risk. The CVSS score is 9.8 regardless of network position. The record confirms exploitation in the wild, and an internal attacker or a compromised internal host can exploit this without authentication.

How we would triage it

This is our judgement.

The commands only read: they print a version or list what is installed. They come from general knowledge of this kind of product, not from the sources, so check them against the vendor's page

before you run them. Words in capitals are for you to fill in.

Step 1: Find where MONAI Label is running

- Check servers used for medical imaging or AI model training for MONAI Label installations and their versions.
- Authentication is off by default; also note whether any instance is reachable from outside the internal network.

```
pip show monailabel
docker images projectmonai/monailabel
```

Step 2: Restrict access while preparing the upgrade

- Block untrusted network access to the MONAI Label server port if upgrading immediately is not possible.
- The vendor's advisory says enabling authentication and limiting the ANNOTATOR role to trusted users reduces risk until the upgrade is applied.

Step 3: Upgrade to 0.8.6

- The release notes name 0.8.6 as the patched version, available as a pip package and a Docker image.
- The release notes also drop Python 3.9 and 3.10 support; 0.8.6 requires Python 3.11 or newer.

Step 4: Review for unexpected file changes if the server was exposed

- The release notes say any 0.8.5 or older server reachable from an untrusted network should be reviewed for unexpected file changes before returning to service.
- The vendor's advisory names SSH authorised keys files, cron directories, and application code as locations an attacker would target.

To close it, read the version from the product itself, not from an old inventory, and save the output with the date. If you cannot update today, write down who owns it, the new date, and what limits exposure until then. CISA's wording adds that you should stop using the product if no mitigation is available.

Send this to your IT provider

Copy it, change the names and dates, and send it. It is a draft for you to check, not a template tested with any provider.

> *We are writing about CVE-2026-133968, a critical path traversal vulnerability in MONAI Label affecting all versions up to and including 0.8.5. The flaw allows an unauthenticated attacker to write arbitrary files to the server, potentially leading to remote code execution. CISA has listed it as actively exploited with a remediation deadline of 26 November 2026. Please confirm whether we run MONAI Label anywhere on our network, and if so, which version and whether it is reachable from outside our internal network. The fixed version is 0.8.6, released 22 November 2026. We need written confirmation of the installed version and whether the upgrade has been applied, by 26 November 2026.*

Signs it was used

An update closes the hole. It does not tell you whether someone used it first. CISA's required action points to its forensic triage requirements (BOD 26-04). Look for:

- Requests to the datastore image and label upload endpoints with path traversal sequences in the image ID or label tag parameters, as the vendor's advisory names these the vulnerable endpoints.
- Files appearing outside the datastore directory, particularly in SSH, cron, or application code locations, with modification timestamps from the period of potential exposure.
- Files in locations the server process can write to with content or timestamps inconsistent with normal deployment, as the vendor's advisory says arbitrary file write is the mechanism of harm.

The weakness behind it

The record classes this as CWE-24 (Path Traversal: '../filedir').

Likely ATT&CK mapping

MITRE ATT&CK is a public catalogue of the steps attackers take, from getting in to causing damage. No source has mapped this flaw to it yet, so the steps below are our prediction from what the record says the flaw allows. They are not MITRE's, the vendor's or CISA's, and an attacker may do other things.

Steps an attacker would likely take with this flaw

The step, in plain words	Stage of an attack	Why we think so

Exploit Public-Facing Application (T1190): Breaking in through a service that faces the internet, such as a web server, VPN or login page.	Initial Access: getting in	The flaw is exploitable without authentication over any network connection; confirmed exploitation in the wild makes initial access via the exposed endpoint likely.
Hijack Execution Flow (T1574): Making a program load the attacker's code instead of the real one.	Persistence: staying in	The advisory names overwriting application code and adding cron entries as ways to achieve execution; writing to those locations would typically hijack the execution flow.

How to use this: ask your IT provider which control or log would show each step, for example whether sign-in attempts from the internet are recorded. If nobody can answer, that is the gap to close first.

MITRE ATT&CK® is a registered trademark of The MITRE Corporation. Technique names are MITRE's.

The sources, in their words

CISA's catalogue entry:

> *Project MONAI MONAI Label contains a path traversal vulnerability that allows an unauthenticated remote attacker to write arbitrary files outside the intended datastore directory, potentially resulting in remote code execution.*

The CVE record, published by Mitre on 20 November 2026:

> *monailabel (aka MONAI Label) before 0.8.6 allows directory traversal, with resultant remote code execution, via the image parameter to the PUT /datastore/image endpoint, as exploited in the wild starting in September 2026.*

Versions of Project-MONIA MONAI Label the record lists as affected:

- all versions before 0.8.6

The affected ranges end before 0.8.6. Those are the first versions the record does not list as affected; confirm them in the vendor's advisory.

CISA's required action:

> *Apply mitigations in accordance with vendor instructions.*

Check your own project

If you build software, you can compare what a project declares with this CVE from a terminal. The tool reads requirements files, package.json, Maven, Gradle, Cargo, CMake and Dockerfile files, runs nothing, and only the CVE id leaves your machine. Naming the CVEs needs no key.

```
lattice-scan project . --cve CVE-2026-133968
```

Each row says in-range, outside-range, undetermined or lead, with the file and line. That is evidence about what a manifest declares, not about what runs, and a person decides. The verdicts, the CSV columns and a pipeline recipe are in the guide.

Check the servers you run

The same id can be checked against a host. Run this on the machine itself, not on the laptop that deploys to it: it reads installed packages (dpkg, rpm or apk), snaps, the hardware (maker, model, firmware), the running container images with their operating system and packages, and the kernel version, compares them on the spot, and only the CVE ids leave the machine; the inventory stays on it, in ~/.lattice-scan. Naming the CVEs needs no key. Do it on each server after a deployment, and again when a new entry lands.

```
lattice-scan --cve CVE-2026-133968
```

A match is a candidate for a person to confirm, not a finding. A distribution can fix a CVE under an unchanged version number, so a match is a candidate; on Debian-family hosts the report also shows whether the package's own changelog names the CVE. By default it opens the images your containers are running and reads their operating system and package list; a deep scan (--deep, a few minutes) opens every local image and the Python, Node, Java and Ruby packages inside them; for what an image declares, use the project check on its Dockerfile.

Guide: namiq.io/lattice/guides/choose-your-channel. MCP server (io.namiq/lattice in the MCP registry): namiq.io/lattice. Package: pypi.org/project/lattice-scan, or namiq.io/apt.

What we do not know yet

- The second path traversal fixed in 0.8.6 (the session upload endpoint) has not been assigned a separate CVE in the sources.
- The sources do not describe how to detect exploitation in MONAI Label's own logs.
- The sources do not say whether a backport fix is available for deployments that cannot move to Python 3.11, which 0.8.6 requires.
- We have not tested this against any environment. Versions and scores are as the sources state them and as at 25 November 2026; the vendor's advisory wins where it differs.

- The sections headed "The bottom line", "Does this apply to you?", "What it is", "Why it matters", "Exposure by situation", "Things people will get wrong", "Likely ATT&CK mapping", "How we would triage it" and "Send this to your IT provider" are our interpretation of the sources, not a statement by CISA or the vendor.

Appendix 1. Severity decoded

CVE-2026-133968: score 9.8, `CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H` As the CVE record gives it. The record's own score can differ from the base score on the NVD record.

Metric	Value	What it means here
Attack vector	Network	Reachable over a network, with no need to be on the same segment.
Attack complexity	Low	No special conditions beyond the flaw itself.
Privileges required	None	No account is needed.
User interaction	None	No one has to do anything.
Scope	Unchanged	Damage stays inside the vulnerable component.
Confidentiality	High	Total loss of confidentiality.
Integrity	High	Total loss of integrity.
Availability	High	Total loss of availability.

Appendix 2. Timeline

Date	Event
20 November 2026	CVE record published by Mitre.
23 November 2026	CISA added it to its list of exploited vulnerabilities.
26 November 2026	CISA's deadline for US federal agencies.

Appendix 3. Exposure by situation

These priorities are our judgement from the attacker prerequisites in the sources. They are not vendor ratings.

Your situation	Priority	Timeframe
MONAI Label 0.8.5 or earlier, reachable from an untrusted network, authentication off (the default)	High	Before 26 November 2026
MONAI Label 0.8.5 or earlier, on an internal network only, authentication disabled	High	This week
MONAI Label 0.8.5 or earlier, authentication on, ANNOTATOR role granted to untrusted users	Medium	This week

MONAI Label 0.8.5 or earlier, authentication on, roles limited to trusted users only	Medium	Next patch cycle
MONAI Label not installed	None	none for this CVE

Appendix 4. Definition of done

- Every Project MONAI MONAI Label you run is listed, with its version.
- Each is outside the affected list, or the vendor's mitigation is in place.
- The version output or a screenshot is saved with the date.
- Your projects' declared dependencies were checked against this CVE, and every row has a written reason.
- Logs for the weeks before the update were searched, and anything odd was written down.
- Anything not done has an owner and a date.

Sources

- CISA Known Exploited Vulnerabilities catalogue: https://www.cisa.gov/known-exploited-vulnerabilities-catalog
- CVE record for CVE-2026-133968: https://www.cve.org/CVERecord?id=CVE-2026-133968
- FIRST: Exploit Prediction Scoring System: https://www.first.org/epss/
- Advisory or release notes read: github.com/Project-MONAI/MONAILabel/security/advisories/GH: https://github.com/Project-MONAI/MONAILabel/security/advisories/GHSA-qc22-rmgc-vg94
- Advisory or release notes read: github.com/Project-MONAI/MONAILabel/releases/tag/0.8.6: https://github.com/Project-MONAI/MONAILabel/releases/tag/0.8.6
- Advisory or release notes read: www.theregister.com/security/2026/11/20/china-attackers-deployed-ai: https://www.theregister.com/security/2026/11/20/china-attackers-deployed-ai-agents-that-autonomously-found-and-exploited-a-0-day-to-take-over-a-university-server/5311382
- Reference in the CVE record: github.com: https://github.com/Project-MONAI/MONAILabel/pull/1915